

Searching Publications on Operating Systems

C.A. Middelburg

Informatics Institute, Faculty of Science, University of Amsterdam,
Science Park 107, 1098 XG Amsterdam, the Netherlands
`C.A.Middelburg@uva.nl`

Abstract. This note concerns a search for publications in which one can find statements that explain the concept of an operating system, reasons for introducing operating systems, a formalization of the concept of an operating system or theory about operating systems based on such a formalization. It reports on the way in which the search has been carried out and the outcome of the search. The outcome includes not only what the search was meant for, but also some added bonuses.

1 Introduction

The study of various issues raised in theoretical work on operating systems and practical work in which operating systems are involved, such as work in the area of digital forensics, calls for an abstract model of an operating system that can be considered an adequate formalization of the concept of an operating system. Moreover, a certain body of theory about operating systems based on the model in question would be helpful in studying the issues concerned. To check whether a more or less appropriate model, as well as theory based on it, already exists, I carried out a search for witnessing publications.

I ran across some specifics of the search by which an ad hoc approach to it would give little confidence in its outcome. This brought me to devise a more or less systematic way to carry out the search. In addition to what was looked for, the search yielded added bonuses. This note reports on the way in which the search has been carried out and the outcome of the search, including the added bonuses.

2 The Approach to Carry out the Search

For the search, Google Scholar and the search engines of relevant journal collections and bibliographies are available. Examples of relevant journal collections are ACM Digital Library, SpringerLink, ScienceDirect, and IEEE Xplore. Examples of relevant bibliographies are ACM Guide, DBLP Computer Science Bibliography and the Collection of Computer Science Bibliographies. The search engines of relevant journal collections and bibliographies have the option to search for publications in which given terms occur in their title, abstract or keywords, but none cover all relevant literature. Google Scholar does not have this

option, i.e. only full text search is possible, but it covers all literature covered by the search engines of relevant journal collections.

In a preparatory phase of the search, we found the following specifics:

- the term “operating system” occurs in extremely many publications whose subject is not operating systems;
- in the early days of operating systems, the term “operating system” is not used in publications on operating systems;
- till 1970, a large part of the publications on operating systems appear in the Communications of the ACM;
- from 1969, a large part of the publications on operating systems appear in ACM SIGOPS Operating Systems Review and the ACM proceedings series SOSP.

By the first two specifics, an ad hoc approach to the search would give little confidence in its outcome. The first specific calls for a search for publications in which the term “operating system” occurs in their title, abstract or keywords, but without special measures this will result in poor coverage of the relevant literature. The second specific complicates a search because not all terms used instead of “operating system” need to be known beforehand, and some terms, such as “supervisor” and “monitor” occur often in different meanings. I took the last two specifics as means to deal with the problems that arise from the first two specifics. This led to the following more or less systematic way to carry out the search:

- first, search for relevant publications by having a look at all publications in all issues of Communications of the ACM till 1970, all issues of ACM SIGOPS Operating Systems Review, and all SOSP proceedings;
- next, search for relevant publications by means of the search engines of relevant journal collections and bibliographies;
- then, search for relevant publications by means of Google Scholar;
- check immediately after finding a relevant publication whether it references directly or indirectly additional relevant publications;
- stop searching by means of Google Scholar when 500 consecutive hits yields no new relevant publication.

The number of publications to be looked at in the first step is about 7,000 (additional publications found by following references not counted). The sum of the numbers of publications on operating systems found by the different search engines of relevant journal collections and bibliographies is about 24,000, but there are many duplicates. I first searched with the search engine of DBLP Computer Science Bibliography, which actually turned out to yield most of the relevant publications found by means of the search engines of relevant journal collections and bibliographies. The number of publications found by a full text search for the term “operating system” with Google Scholar is about 500,000. I looked only at the first 1,000 hits, because Google Scholar yielded very few new relevant publications.

3 A Summary of the Main Outcome of the Search

I searched for publications in which one can find statements that explain the concept of an operating system, reasons for introducing operating systems, a formalization of the concept of an operating system or theory about operating systems based on such a formalization. It turned out that the number of such publications is very small.

In [6], Codd and others give motivation for, requirements for, and functions of a multiprogramming operating system. This can be taken for a preparation to the formulation of the scheduling problem in multiprogramming operating systems in [4] and the description of a scheduling algorithm for a multiprogramming operating system in [5]. Probably these three articles belong to the first important articles on operating systems. It is often stated that Strachey's article on multiprogramming operating systems [24] is the first important article on multiprogramming operating systems, but it appears that a more moderate statement is more appropriate.¹

Apart from the attempt of Codd and others in [6], few serious attempts have been made to explain the concept of an operating system. Dennis and Van Horn make a serious attempt in [14] and Denning makes another serious attempt in [13], but most other attempts cannot be called serious. Examples of non-serious attempts are one-liners like "an operating system is an extended machine and a resource manager" and enumerations of the usual terms for the basic constituents of an operating system.

Apart from the reasons given by Codd and others in [6], reasons for introducing operating systems are seldom given. Cloot gives good reasons in [3], an article whose sole aim is to explain why the need for operating systems has arisen, but usually the reasons are not more advanced than "it is useful to have an operating system available".

In [25], Yates and others give an abstract model of an operating system, using input/output automata, which looks to be a model that can be considered an adequate formalization of the concept of an operating system.² Apart from this, publications in which abstract models of an operating system are given that can be considered an adequate formalization of the concept of an operating system are virtually absent. In [17], Heistand gives an abstract model of a part a specific time-sharing operating system, using finite-state automata. It is questionable whether this model can be considered an adequate formalization of the concept of an operating system. In [11], Degtyarev and Kalinichenko give a very abstract model of an operating system. Because it abstracts also from essentials of an operating system, this model cannot be considered an adequate formalization of the concept of an operating system. Neither the queueing-theoretical models of process scheduling in time-sharing operating systems, such as the models of Kleinrock and Coffman [19,7], are the models we have in mind.

¹ Strachey's paper can only be obtained by ordering a hard copy at the National Archive of the United Kingdom.

² The paper of Yates and others actually gives two models. The abstract model is the model that is called the user level model in the paper.

In [21], McKenney mentions that a conceptual model of an operating system based upon analysis of operating systems and a literature study is given in his Ph.D. thesis. Little more is known about this model: the thesis is not cited in other publications and whether it can still be obtained is doubtful. In any case, it is unlikely that the model of McKenney is not superseded by the model of Yates and others mentioned in the previous paragraph.

Publications on theory about operating systems themselves are totally absent. In publications on operating systems that are of a theoretical nature, one finds only theory about internals of operating systems such as process scheduling and resource allocation. For early articles presenting such theory, see e.g. [19,7] and [10,12], respectively.

From this outcome of the search, we conclude that there exists only one abstract model of an operating system that can be considered a more or less appropriate formalization of the concept of an operating system, viz. the abstract model given by Yates and others in [25], and that there does not exist theory based on that model. Moreover, we conclude that the operating systems community pays little attention to clarifying adequately what is an operating system and giving motives for introducing operating systems. All this raises the question what most publications on operating systems are about. It turns out that they mainly concern the following:

- principles of operating system design;
- theory and techniques related to internals of operating systems such as process scheduling and resource allocation;
- issues concerning operating systems for multi-processor computers and operating systems for networks of distributed computers;
- operating system support for security, privacy, fault-tolerance, multi-media applications, etc.;
- designs of, analyses of, and experiences with specific operating systems.

It is striking that most of these publications give little insight in the concept of an operating system. My findings of the search agree with the finding of the study of courses and textbooks presented in [9].

4 The Added Bonuses of the Search

I found many statements that try to explain the concept of an operating system in ways that lack generality, abstractness or preciseness. However, I could easily distill the essence from them. This led to the following explanation of the concept of an operating system.

An operating system is a system that provides a convenient execution environment for programs that allows for multiple programs with shared resources to be executed concurrently. An operating system is responsible for:

1. loading programs and starting their execution;
2. scheduling the programs in execution;

3. allocating resources to the programs in execution;
4. preventing interference between the programs in execution;
5. controlling the use of main memory by the programs in execution;
6. storing and retrieving data organized into files and directories on secondary storage devices;
7. receiving data from input devices and sending data to output devices;
8. communicating data over computer networks;
9. controlling peripheral devices.

It is customary to distinguish the following basic constituents in an operating system:

- process management, responsible for 1, 2, 3 and 4;
- memory management, responsible for 5;
- file management, responsible for 6;
- input/output management, responsible for 7;
- network management, responsible for 8;
- device management, responsible for 9.

Process management and a part of memory management are needed to provide an execution environment for programs that allows for multiple programs with shared resources to be executed concurrently. Device management, network management, input/output management, file management, and a part of memory management are needed to provide a *convenient* execution environment, because they hide interrupts, networking protocols, device-dependent input, output and storage, physical memory size, etc.

Operating systems can be classified as:

- single-user or multi-user;
- non-interactive or interactive;
- single-tasking, non-preemptive multi-tasking or preemptive multi-tasking.

Actually, the explanation given above is an explanation of the concept of an multi-tasking operating system. Single-tasking operating systems are border cases of operating systems: the maximal number of programs that can be executed concurrently is only one. Clearly, a multi-tasking operating system is a more general concept than a single-tasking operating system.

Batch operating systems are multi-user, non-interactive, single-tasking operating systems. Multiprogramming operating systems are multi-user, non-interactive, (non-preemptive or preemptive) multi-tasking operating systems. Time-sharing operating systems are multi-user, interactive, preemptive multi-tasking operating systems. Ryckman mentions in [23] that the Input/Output System for the IBM 704 computer, which was developed by General Motors and North American Aviation and became operational in 1956, is the first batch operating system. The Atlas Supervisor of Kilburn and others [18], which was developed over the period 1957–1961, is generally considered the first multiprogramming operating system. The experimental time-sharing operating system CTSS of Corbató and others [8], which was developed over the period 1961–1963, is generally considered the first time-sharing operating system.

I accidentally found additional interesting facts about the history of operating systems which seem to be overlooked in all historical overviews.

We already mentioned that Strachey’s article on multiprogramming operating systems [24] is often mentioned as the first important article on multiprogramming operating systems, but that there exists an article by Codd and others, namely [6], which probably belongs to the first important articles on multiprogramming operating systems as well. In the article of Codd and others is referred to [2], which is an article on a program interruption system published in 1957. I found that the latter article already mentions multiprogramming without explanation. Therefore, I decided to search for earlier publications in which the term “multiprogramming” occur. I found that multiprogramming is mentioned as a programming technique for sorting by merging in a paper by Friend on sorting [16] published in 1956. Moreover, multiprogramming is shortly explained by means of an example in a paper by Rochester [22] published in 1955. To my knowledge, these papers are never mentioned in historical remarks made in the literature in question.

The idea of a time-sharing operating system is usually attributed to McCarthy, who proposed to develop a time-sharing operating system at MIT in 1959. His recollections about this matter can be found in [20]. The term “time-sharing” has caused confusion about the origin of the idea of a time-sharing operating system, because this term is used before 1959 by Everett and others in [15] and Bemmer in [1] for preemptive multi-tasking in a non-interactive setting. In [24], Strachey uses the term “time-sharing” for “multi-tasking” as well. On that account, Corbató and others [8] attributes the idea of a time-sharing operating system incorrectly to Strachey.

5 Concluding remarks

To check whether an abstract model of an operating system that can be considered a more or less appropriate formalization of the concept of an operating system and a certain body of theory about operating systems based on it already exist, I have carried out a search for witnessing publications. To be able to judge the suitability of existing models, I have also looked for publications that contain statements explaining the concept of an operating system and publications that give reasons for introducing operating systems.

From the outcome of the search, I conclude that there exists only one more or less appropriate model, viz. the abstract model given by Yates and others in [25], and that there does not exist theory based on that model. In addition, I conclude that the operating systems community pays little attention to clarifying adequately what is an operating system and giving motives for introducing operating systems. However, I found many statements that try to explain the concept of an operating system in ways that lack generality, abstractness or preciseness. By distilling the essence from those statements, I have obtained an explanation of the concept of an operating system that could serve as the starting-point of the development of an abstract model of an operating system. It happens that

the above-mentioned operating system model of Yates and others agrees with this explanation.

References

1. Bemert, R.W.: How to consider a computer. *Automatic Control* pp. 66–69 (1957)
2. Brooks Jr., F.P.: A program-controlled program interruption system. In: *IRE-ACM-AIEE '57 (Eastern)*, pp. 128–132. ACM Press (1958)
3. Clood, P.L.: What is the use of operating systems? *Computer Journal* **7**(4), 249–254 (1965)
4. Codd, E.F.: Multiprogram scheduling: Parts 1 and 2. introduction and theory. *Communications of the ACM* **3**(6), 347–350 (1960)
5. Codd, E.F.: Multiprogram scheduling: Parts 3 and 4. scheduling algorithm and external constraints. *Communications of the ACM* **3**(7), 413–418 (1960)
6. Codd, E.F., Lowry, E.S., McDonough, E., Scalzi, C.A.: Multiprogramming STRETCH: Feasibility considerations. *Communications of the ACM* **2**(11), 13–17 (1959)
7. Coffman, E.G., Kleinrock, L.: Feedback queueing models for time-shared systems. *Journal of the ACM* **15**(4), 549–576 (1968)
8. Corbató, F.J., Merwin-Daggett, M., Daley, R.C.: An experimental time-sharing system. In: *AIEE-IRE '62 (Spring)*, pp. 335–344. ACM Press (1962)
9. Creak, G.A., Sheehan, R.: A top-down operating systems course. *ACM SIGOPS Operating Systems Review* **34**(3), 69–80 (2000)
10. Dahm, D.M., Gerbstadt, F.H., Pacelli, M.M.: A system organization for resource allocation. *Communications of the ACM* **10**(12), 772–779 (1967)
11. Degtyarev, E.K., Kalinichenko, S.P.: Implementation of levels of abstraction in an operating system. *Cybernetics and Systems Analysis* **16**(2), 253–260 (1980)
12. Denning, P.J.: The working set model for program behavior. *Communications of the ACM* **11**(5), 323–333 (1968)
13. Denning, P.J.: Third generation computer systems. *ACM Computing Surveys* **3**(4), 175–216 (1971)
14. Dennis, J.B., Van Horn, E.C.: Programming semantics for multiprogrammed computations. *Communications of the ACM* **9**(3), 143–155 (1966)
15. Everett, R.R., Zraket, C.A., Benington, H.D.: SAGE: A data-processing system for air defense. In: *IRE-ACM-AIEE '57 (Eastern)*, pp. 148–155. ACM Press (1957)
16. Friend, E.H.: Sorting on electronic computer systems. *Journal of the ACM* **3**(3), 134–168 (1956)
17. Heistand, R.E.: An executive system implemented as a finite-state automaton. *Communications of the ACM* **7**(11), 669–677 (1964)
18. Kilburn, T., Payne, R.B., Howarth, D.J.: The Atlas supervisor. In: *AFIPS '61 (Eastern)*, pp. 279–294. ACM Press (1961)
19. Kleinrock, L.: Time-shared systems: A theoretical treatment. *Journal of the ACM* **14**(2), 242–261 (1967)
20. Lee, J.A.N., McCarthy, J., Licklider, J.C.R.: The beginnings at MIT. *IEEE Annals of the History of Computing* **14**(1), 18–30 (1992)
21. McKenney, J.L.: Simultaneous processing of jobs on an electronic computer. *Management Science* **8**(3), 344–354 (1962)
22. Rochester, N.: The computer and its peripheral equipment. In: *AIEE-IRE '55 (Eastern)*, pp. 64–69. ACM Press (1955)

23. Ryckman, G.F.: The IBM 701 computer at the General Motors Research Laboratories. *IEEE Annals of the History of Computing* **5**(2), 210–212 (1983)
24. Strachey, C.: Time sharing in large fast computers. In: *International Conference on Information Processing*, pp. 336–341. UNESCO (1959)
25. Yates, D., Lynch, N., Seltzer, M., Luchangco, V.: I/O automaton model of operating system primitives (1999). Bachelors thesis HU 92.99, Harvard University